Delay-Tolerant Networking

June 5, 2003

Scott Burleigh
DTN Research Group, JPL
818-393-3353
Scott.Burleigh@jpl.nasa.gov

Scenario

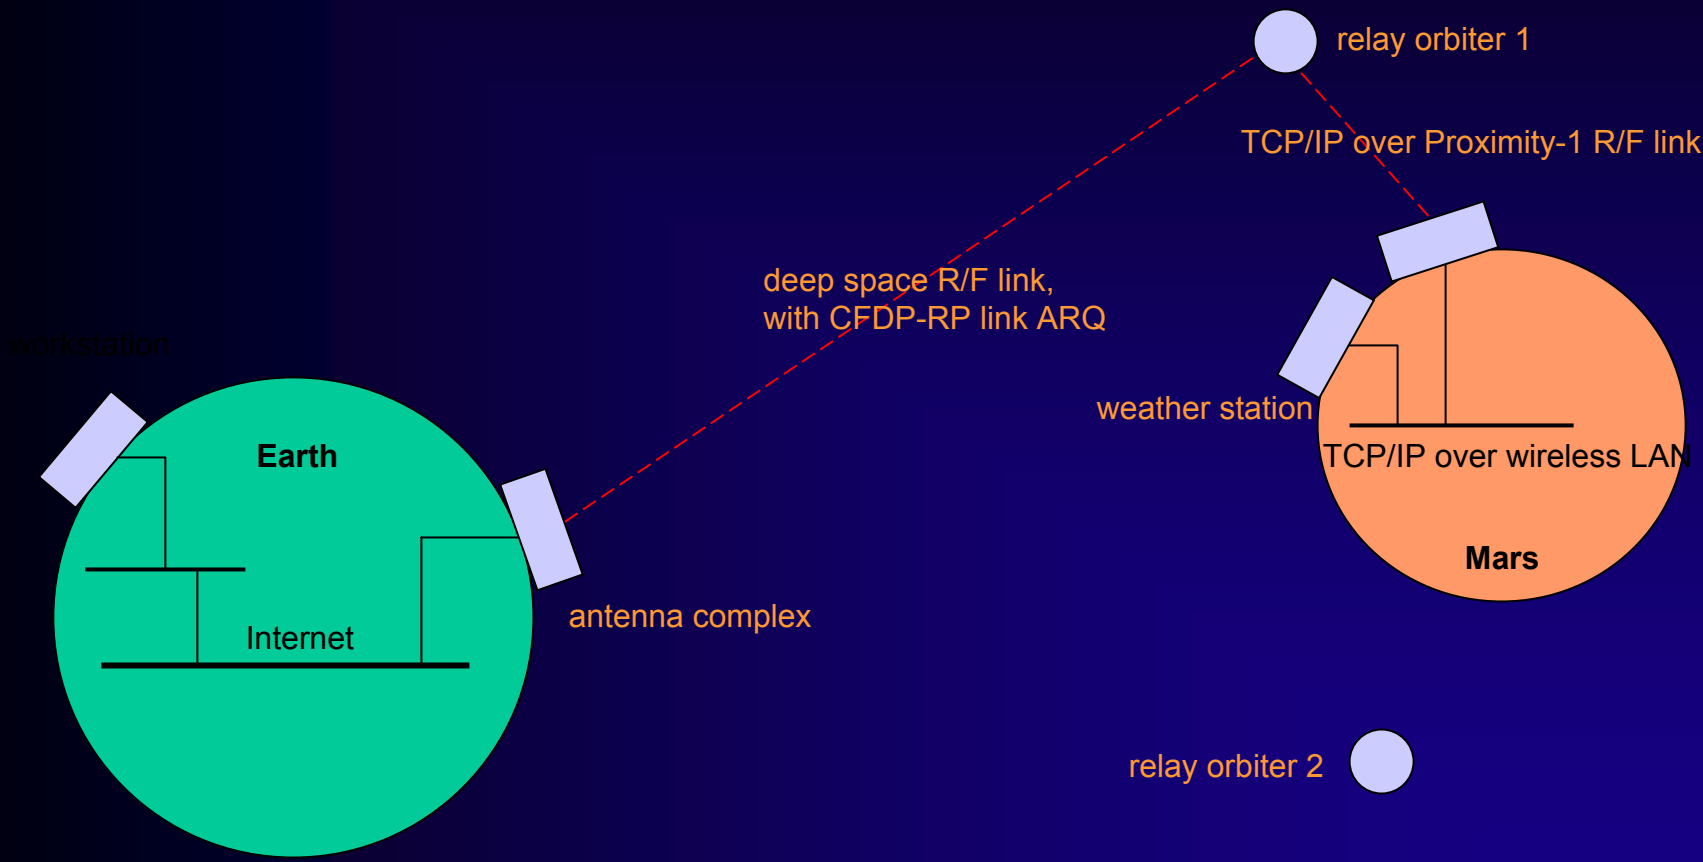
relay orbiter 1
TCP/IP over Proximity-1 R/F link
deep space R/F link,
with CFDP-RP link ARQ
workstation
weather station
TCP/IP over wireless LAN
Earth
Mars
Internet
antenna complex
relay orbiter 2

Why Not IP End-to-End?

- Problems with TCP:
  - Connection time (one round trip) may exceed duration of communication opportunity.
  - In-order delivery means data loss delays data arrival by at least one round trip.
  - Long round-trip times retard recovery from data loss (interpreted as congestion, reducing data rate).
- Problem with end-to-end ARQ (either TCP or application-layer ARQ over UDP): end-to-end retransmission requires original sender to retain retransmission buffer for an e-to-e round trip.

Why Not IP End-to-End? (cont'd)

- Problems with routing protocols:
  - BGP uses TCP, performs poorly when TCP is unable to keep a connection established.
  - Route computation is based on probes and timeouts; loss of connectivity can result in premature timeout, thus a connectivity false negative.
  - Transient partitioning due to scheduled intermittent connectivity may be interpreted as loss of connectivity to the destination – no route can be computed at all.

Alternative: Delay-Tolerant Networking

- Use protocols at all layers of stack that are best suited to each environment.
- Above them, provide an overlay network protocol that applications can use end-to-end.
- No end-to-end expectation of:
  - continuous connectivity; low or constant transmission latency; low error rate or low congestion
  - high transmission rate or symmetrical data rates
  - common name or address expression syntax or semantics; data arrival in transmission order

DTN Principles

- Postal model of communications.
  - Abandon telephonic model. Don't design for interactive conversation. When submitting a request, "bundle" with it the answers to all possible questions.
  - Overlay protocol is named "Bundling".
- Tiered functionality: rely on underlying "regional" protocols as heavily as possible, do the rest in Bundling.
- Terseness.

Tiered Forwarding

- Regional network protocols (e.g., IP) do local forwarding.
- Bundling does end-to-end forwarding across region boundaries. *Deferred transmission*.
- Bundle (message) source and destination IDs must include:
  - Region ID (meaningful to Bundling)
  - Regional endpoint ID (meaningful to the regional network protocol)
- Region IDs function as addresses.
- Regional endpoint IDs are names that are *late bound* to regional addresses upon arrival at the destination region.

Other Tiered Functionality

- Tiered routing:
  - Regional routing protocols
  - Contact-sensitive bundle routing
- Tiered ARQ:
  - Regional ARQ (e.g., TCP, CFDP's retransmission procedures)
  - Bundle-layer ARQ: custody transfer
- Tiered security:
  - Hop-by-hop bundle agent authentication at Bundle layer to protect infrastructure
  - End-to-end confidentiality, integrity at application layer

Other Tiered Functionality (cont'd)

- Tiered congestion avoidance:
  - Regional protocols deal with regional congestion.
  - Bundle layer detects Bundling congestion, respond to it by invoking (tiered) flow control.
- Tiered flow control:
  - Regional flow control may be protocol-based (Internet) or managed, rate-based (deep space).
  - Inhibition of custody acceptance at Bundle layer.
- Tiered coding:
  - Regional coding:
    - Bundle header compression in adapters to regional protocols.
    - Other coding as needed is performed by regional protocols, possibly at multiple layers of stack.
  - Optional erasure coding at Bundle layer.

Other Tiered Functionality (cont'd)

- Tiered fragmentation and reassembly:
  - Bundling fragments bundles from awareness of contact duration.
    - Proactive for scheduled or predicted contacts.
    - Reactive for opportunistic contacts.
  - Regional protocols do further fragmentation from awareness of (e.g.) MTU size.

Other DTN Functionality

- Resilient delivery: destination service agent may not be running at the time a bundle destined for it arrives.
  - *Deferred delivery*: wait until destination starts.
  - *Reanimation*: start the destination, then deliver bundle to it.
- Postal service levels:
  - Priority levels: low, standard, high
  - Service notifications:
    - Notice of initial transmission, i.e., notice of mailing
    - Notice of delivery to the ultimate destination application , i.e., return receipt
    - Report of route taken, i.e., delivery record

Example of Bundling Data Flow

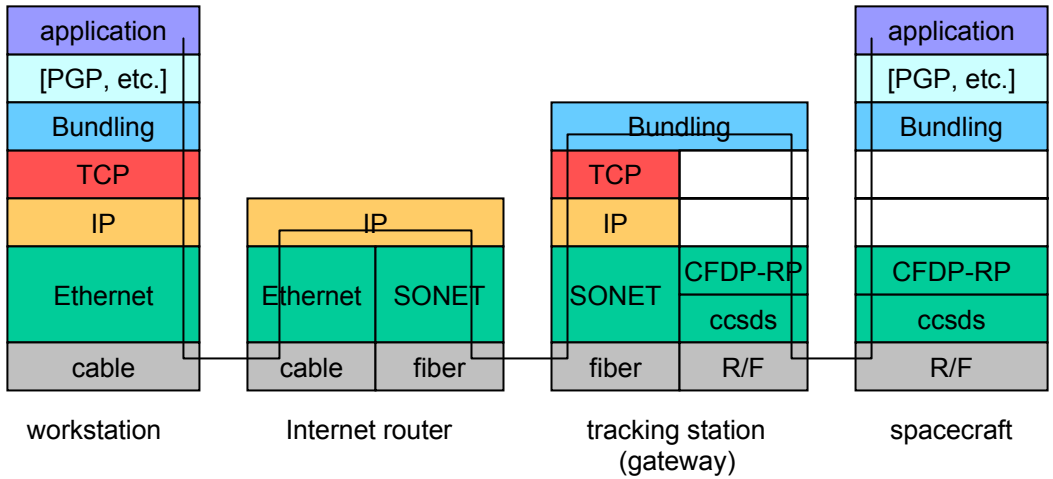

Example of Tunneling/RLI Data Flow

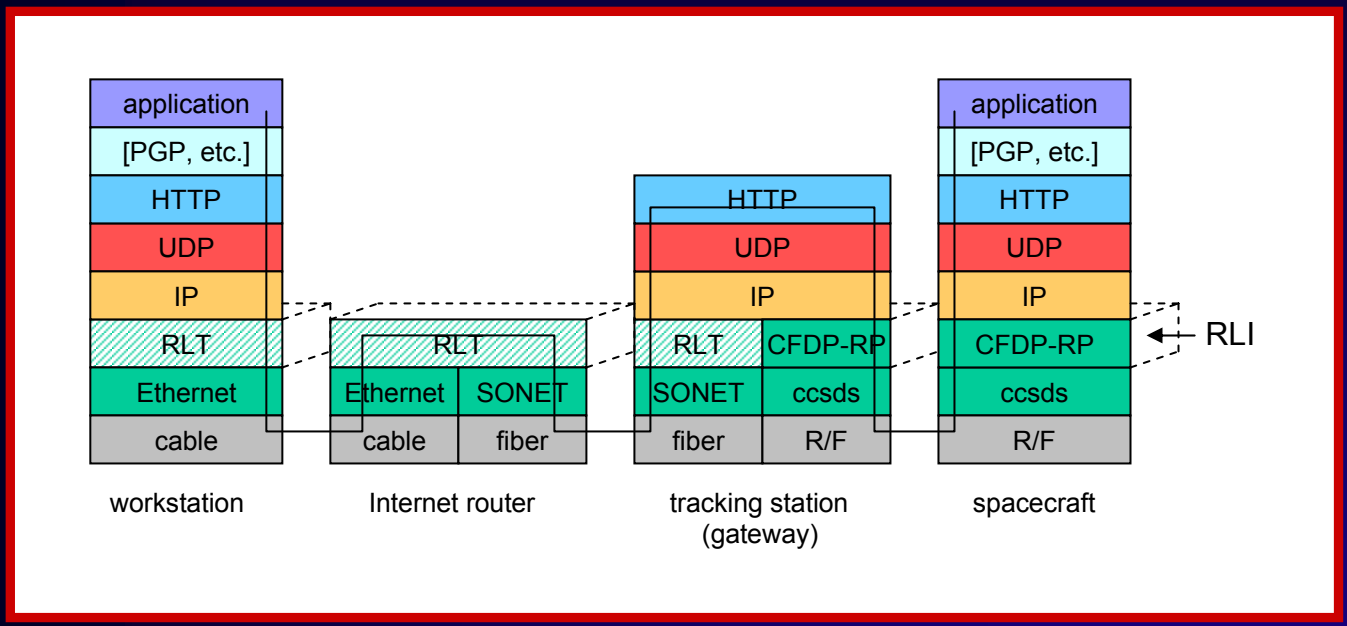

Summary

- Emerging network configuration problems are difficult to handle by simply extending the Internet.
- Delay-Tolerant Networking generalizes the Internet architecture to address these problems in a simple, robust way.

• NASA DATA SYSTEM STANDARDS PROGRAM •

1. Interplanetary Internet: An Architectural Framework for Space Internetworking: Adrian Hooke
2. User Data Services for Internet Based Spacecraft Applications: Joe Smith
3. CCSDS File Delivery Protocol (CFDP): Tim Ray
4. Internet Protocol Based Standards for Spacecraft Onboard Interfaces: Joe Smith
5. Standard Spacecraft Interfaces and IP Network Architectures: Jane Marquart
6. Standard Transport and Network Capabilities: Bob Durst
7. Next Generation Space Internet: Standards and Implementation: Keith Scott
8. Secure Space Networking: Howie Weiss
9. Delay Tolerant Networking: Scott Burleigh
10. CCSDS Link Layer Protocol Suite: Greg Kazz